

Anomaly-Sensitive Dictionary Learning for Unsupervised Diagnostics of Solid Media

Jeffrey M. Druce^a, Jarvis D. Haupt^b, Stefano Gonella^{a,*}

^aDepartment of Civil Engineering, University of Minnesota, Minneapolis, MN, USA

^bDepartment of Electrical and Computer Engineering, University of Minnesota, Minneapolis, MN, USA

Abstract

This paper proposes a strategy for the detection and triangulation of structural anomalies in solid media. The method revolves around the construction of sparse representations of the medium's dynamic response, obtained by learning instructive dictionaries which form a suitable basis for the response data. The resulting sparse coding problem is recast as a modified dictionary learning task with additional spatial sparsity constraints enforced on the atoms of the learned dictionaries, which provides them with a prescribed spatial topology that is designed to unveil anomalous regions in the physical domain. The proposed methodology is model agnostic, i.e., it forsakes the need for a physical model and requires virtually no *a priori* knowledge of the structure's material properties, as all the inferences are exclusively informed by the data through the layers of information that are available in the intrinsic salient structure of the material's dynamic response. This characteristic makes the approach powerful for anomaly identification in systems with unknown or heterogeneous property distribution, for which a model is unsuitable or unreliable. The method is validated using both synthetically generated data and experimental data acquired using a scanning laser Doppler vibrometer.

Keywords: Dictionary learning, Sparse coding, Anomaly detection, Wave propagation, Modal analysis, Structural diagnostics, Laser vibrometer

I. Introduction

The past three decades have witnessed the advent of a variety of structural diagnostics methodologies based on guided waves [1, 2]. The detection and triangulation principle underlying all methods based on guided waves (the pulse-echo or pitch-catch approaches [3]) can be seen as a structural analog of the *radar* problem: waves are generated and received by transmitter-receiver pairs distributed over the test specimen, a signature of wave scattering is captured along each transmitter-receiver path, and the position of the defect is subsequently triangulated using data from multiple transducers. Numerous efforts have been dedicated to the construction of damage location estimators from measurements acquired by sparse arrays of sensors. Popular approaches include: statistical methods [4], acoustic imaging techniques [5], singular value decomposition [6, 7], spatial optimization of the sensor networks [8], and methods based on the time reversal operator [9, 10].

Pitch-catch methods allow anomaly triangulation using parsimonious sensors data, which makes them ideally suited for online or *in-situ* SHM applications [11, 12], where it is crucial that the acquisition system

*Corresponding Author

Email address: sgonella@umn.edu (Stefano Gonella)

is highly portable and easily deployable. On the other hand, they, like all methods based on the radar paradigm, can suffer major weaknesses when the assumptions on the ideality of the medium are relaxed - a common scenario in the context of aging materials and damage formation. These methods rely in fact on the possibility to detect the individual scattered signals, estimate with some precision the associated times of flight, and finally triangulate the scattering sources by applying some direct knowledge of the medium properties (e.g., the wave speed). These tasks are often hard to accomplish in the case of highly heterogeneous media, materials with extreme internal complexity (e.g., random microstructures) or materials experiencing severe property degradation, for which a material model is either unknown or unreliable.

Recently, the field of structural diagnostics has been flooded with methodologies originally developed within the field of machine learning (ML) [13]. The picture that emerges is one of a fast-growing field that is quickly incorporating inputs from parallel disciplines. Examples of ML techniques used in support of wave-based diagnostics include neural networks [14], matching pursuit decomposition [15–18], support vector machines [19] and compressive sensing [20]. It is worth pointing out that the majority of the existing ML approaches for diagnostics involve supervised learning techniques for feature classification, whose use may be hindered by the need for large training data sets and databases.

In parallel, another powerful class of diagnostic methodologies has stemmed from the availability of laser-based acquisition systems [21, 22]. By means of a Scanning Laser Doppler Vibrometer (SLDV) it is possible to perform non-contact measurements of the velocity of points belonging to a (potentially very dense) scanning grid defined on an object’s surface, which enables full spatial reconstruction of its vibration or wave response. A number of dedicated image processing techniques have been developed in conjunction with laser experiments to meet desired identification and visualization criteria; among them are methods based on space-time DFT [23, 24], wavenumber-space filters [25] and Laplace filters [26, 27]. In sheer contrast with the radar triangulation approach, laser-based diagnostics promote a different paradigm where the inference is performed directly on a data-rich, spatially reconstructed response. While the acquisition of richer data poses additional challenges in terms of sensing system requirements, it opens new avenues for inference strategies with superior accuracy and robustness.

The possibility to turn the attention from analyzing sparse time history data to operating upon spatially reconstructed waveforms, coupled with the availability of algorithms to mine complex data structures, represents a significant shift in perspective for the diagnostics problem. At the core of this new approach is the notion that, from a data standpoint, a wavefield is essentially a data cube, slices of which represent snapshots of the dynamic response at different time instants. Treating an evolving wavefield as a collection of images immediately presents the opportunity to revisit the anomaly detection problem with the mind-set and methodology of image processing and computer vision (CV) [28–31]. For example, the problem of detecting anomalies in the physical medium has a data equivalent in the problem of identifying atypical patterns in the data structure [32, 33]. One of the essential concepts behind this approach is the notion that, in every region of the domain that is sufficiently far from a defect, the displacement time histories will exhibit some similar, but unknown, “typical” behavior, while the time histories recorded in spatial regions in the immediate vicinity of a defect will exhibit some (also unknown) signature of the defect that is different from the typical response observed in the bulk of the domain. The regions exhibiting atypical behavior are referred to as *salient*. When only a few regions exhibit atypical behavior, the notion of *salience* can be viewed as a generalization of the concept of *sparsity*, which has played a central role in signal processing,

statistics, and machine learning research in recent years; see, e.g., [34–41]. In the context of wave-based structural diagnostics, an approach based on notions of saliency has been recently investigated in [42].

The concept of sparsity is also at the core of recent efforts in *dictionary learning* methods (e.g., [43–47]), whose objective is to seek a sparse decomposition of a data set in terms of a basis (typically, but not necessarily, overcomplete) that is learned directly from the data. The idea of representing a signal using a small set of atoms of a learned dictionary instead of prescribed basis functions, such as sinusoids or wavelets, has shown to be a powerful and versatile tool in many signal and image processing applications, including image denoising and restoration, texture synthesis and classification [48–50], audio classification [51] and source separation [52], and medical imaging [53].

In this paper, we introduce a sparse coding approach to the structural diagnostics task based on the reinterpretation of the dictionary learning problem as a generalized, anomaly-sensitive form of modal analysis. We show that we can steer the outcome of the sparse coding problem toward the identification of anomalous features in the medium through the introduction of a data model whose parameters embody features associated with certain morphological, structural and/or behavioral characteristics that we presume should be exhibited in the response data. The resulting inference problem is fully model-agnostic and baseline-free, in that it does not require any *a priori* knowledge of the material model of the medium (e.g., governing equations, material properties); therefore our proposed approach is well suited to study media whose material model is unknown, due for example to heterogeneity in the property distribution, or unreliable, as in the presence of material degradation extended over large portions of the domain. The approach is also database-free, as the construction of the data model is unsupervised, i.e., does not involve any training sets.

The remainder of the paper is organized as follows. In Section II we present the dictionary learning paradigm as a natural (and versatile) generalization of conventional modal analysis. In Section III we show how the sparse coding problem can be modified to capture localized features in the response data. The effectiveness of the method is demonstrated in Section IV using synthetic data as well as data acquired using a scanning laser Doppler vibrometer. In Section V we provide some final remarks and we discuss a few possible directions for future work.

II. Dictionary Learning as Generalized Modal Analysis

In the context of structural dynamics, it is well known that the arbitrary motion of an undamped continuous system can be described as a linear combination of its natural modes, i.e., the dynamic displacement field $w(\mathbf{x}, t)$ can be expressed as

$$w(\mathbf{x}, t) = \sum_{j=1}^{\infty} \gamma_j \phi_j(\mathbf{x}) e^{i\omega_j t} \quad (1)$$

where $\phi_j(\mathbf{x})$ are the mode shapes, which are functions of the position vector \mathbf{x} , ω_j are the natural frequencies of the system, and γ_j are modal amplitudes that are determined from the initial conditions of the problem. If the summation in Eq. 1 is truncated to M terms, it yields an approximation of the displacement field. If we consider a discrete system with N degrees of freedom (e.g., a lumped-parameter system or a continuous system upon spatial discretization), whose response is sampled at T discrete time instants, we can update




Figure 1: Matrix representation of conventional modal decomposition. Shaded columns of Φ represent spatial functions associated with the digitized time functions appearing as shaded rows of Λ .

Eq. 1 and write the motion at a time instant t_k as

$$\mathbf{w}(t_k) = \sum_{j=1}^N \gamma_j \phi_j e^{i\omega_j t_k} \quad k = 1, \dots, T. \quad (2)$$

where \mathbf{w} is an $N \times 1$ array of degrees of freedom and ϕ_j 's are $N \times 1$ arrays representing the vectorized discrete mode shapes. We note that the exact response aggregates N modal contributions, i.e., the number of modes equals the number of degrees of freedom. The modal decomposition is often recast in the form of a matrix multiplication as $\mathbf{X} = \Phi \Lambda$, where $\mathbf{X} \in \mathbb{R}^{N \times T}$ is the matrix of the response data, containing a length T discrete time history for each of the N degrees of freedom, $\Phi \in \mathbb{R}^{N \times N}$ is the (square) modal matrix whose j^{th} column corresponds to the vectorized mode shape ϕ_j , $\Lambda \in \mathbb{R}^{N \times T}$ is a matrix whose λ_j row corresponds to the time-harmonic evolution of the j^{th} modal coordinate, and the modal amplitudes γ_j are assumed absorbed into the rows of Λ . In short, the modal decomposition can be expressed as a representation in terms of a purely spatial component (a matrix of discrete mode shapes) and a purely temporal one (a matrix of discrete harmonic functions). This decomposition is schematically illustrated in Fig. 1.

As a generalization of this approach, we may wish to construct an alternative approximate representation of the response which is exclusively *learned* from the response data (without the need for direct knowledge of the system's mechanical properties) but, unlike modal decomposition, relaxes *a priori* assumptions on the mathematical form of the basis functions. In analogy with the description in terms of modes, we seek a representation that still involves a pair of spatial and temporal matrices, but in which the orthogonality constraint on the columns of the spatial matrix has been relaxed. In compact form, we write

$$\mathbf{X} \approx \mathbf{D} \mathbf{A}, \quad (3)$$

where the matrix \mathbf{D} is a *dictionary* matrix, whose columns (called *atoms*) form a set of spatial basis functions, and \mathbf{A} is a corresponding coefficient matrix, whose rows may still be viewed as digitized functions of time. In this representation, the columns of \mathbf{D} play a role that is functionally analogous to the mode shapes ϕ_j , therefore we refer to the atoms of \mathbf{D} as “*pseudo-mode*” shapes of the system.

Factorizations of the form of Eq.(3) are, of course, not unique without further qualifications; indeed, even the discretized form of standard modal decomposition may be written using this formalism. In what

follows, we adopt this type of general factorization model in order to facilitate “tuning” of our learned data representations, so that they become inherently sensitive to wave propagation characteristics that are indicative of material defects. We can encapsulate such characteristics in terms of the structure(s) that we *prescribe to* or *enforce upon* the learned factors \mathbf{D} and \mathbf{A} . Suppose that we specify classes of candidate dictionaries and candidate coefficient matrices, denoted by \mathcal{D} and \mathcal{A} respectively, so that each $\mathbf{D} \in \mathcal{D}$ and each $\mathbf{A} \in \mathcal{A}$ exhibits structural characteristics that we wish to impose on the dictionary atoms and their coefficients. Then, given \mathbf{X} , the aim of the representation task becomes to find specific factors $\hat{\mathbf{D}} \in \mathcal{D}$ and $\hat{\mathbf{A}} \in \mathcal{A}$, such that $\mathbf{X} \approx \hat{\mathbf{D}}\hat{\mathbf{A}}$. Formally, our approach will be to identify these factors by solving a (constrained) version of a least-squares problem of the form

$$\{\hat{\mathbf{D}}, \hat{\mathbf{A}}\} = \arg \min_{\mathbf{D} \in \mathcal{D}, \mathbf{A} \in \mathcal{A}} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{D}\mathbf{a}_t\|_2^2, \quad (4)$$

where \mathbf{x}_t and \mathbf{a}_t denote the t -th columns of \mathbf{X} and \mathbf{A} , respectively, and the notation $\|\cdot\|_2^2$ denotes the squared ℓ_2 , or Euclidean, norm.

This line of thinking is motivated by recent efforts in *dictionary learning* [43–47] whose objectives are factorizations characterized by dictionaries that may be overcomplete (having more columns than rows) with corresponding coefficient matrices that are *sparse* (having relatively few nonzero entries). In terms of (4) above, dictionary learning tasks may be described as optimizations over a set \mathcal{D} of matrices having N rows and some user-specified number of columns (say K), and a corresponding set \mathcal{A} of $K \times T$ coefficient matrices having no more than $s < K$ non zeros per column. Enforcing sparsity on the columns of \mathbf{A} may be accomplished by imposing a set of constraints of the form $\|\mathbf{a}_t\|_0 \leq s$ for all $t = 1, 2, \dots, T$, on the elements $\mathbf{A} \in \mathcal{A}$, where the notation $\|\mathbf{a}_t\|_0$ denotes the ℓ_0 or counting norm of \mathbf{a}_t , which essentially measures how many of its entries are nonzero¹, and $s > 0$ is a specified sparsity level.

Optimizations of this form (having ℓ_0 constraints) are well-known to be combinatorial in nature. Thus, modern dictionary learning efforts either resort to greedy methods [46], or attempt to *relax* the ℓ_0 constraint on each column to yield computationally tractable constraints. Often, this entails replacing the ℓ_0 norm arising in the constraints on the columns of \mathbf{A} by its closest *convex* surrogate [54]. In terms of the terminology introduced for our problem, this latter approach would prescribe casting the dictionary learning task in terms of an optimization of the form

$$\{\hat{\mathbf{D}}, \hat{\mathbf{A}}\} = \arg \min_{\mathbf{D} \in \mathcal{D}, \mathbf{A} \in \mathcal{A}} \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{D}\mathbf{a}_t\|_2^2 + \lambda \|\mathbf{a}_t\|_1, \quad (5)$$

where \mathcal{A} is the set of all $K \times T$ matrices, \mathcal{D} is the set of $N \times K$ matrices satisfying $\|\mathbf{d}_j\|_2 \leq 1$ for each column $j = 1, 2, \dots, K$, and $\|\mathbf{a}_t\|_1 \triangleq \sum_{i=1}^K |A_{i,t}|$ denotes the ℓ_1 norm of \mathbf{a}_t . Here, $\lambda > 0$ is a (user-specified) regularization parameter that trades off a global “goodness of fit” of the approximation, as quantified by the ℓ_2 term in the objective function, with the ℓ_1 term on the columns of \mathbf{A} ; larger values of λ tend to result in sparser \mathbf{A} having fewer non zeros per column. We now propose an extension of the (sparse) dictionary

¹Strictly speaking, the function $\|\cdot\|_0$ does not satisfy all of the required characteristics for it to be a proper norm. In particular, it fails the homogeneity property, in that for a vector \mathbf{a} we do not have $\|c\mathbf{a}\|_0 = |c| \|\mathbf{a}\|_0$ for all constants c . Nevertheless, it has become common in the sparse inference literature to use the norm descriptor for this function; we adopt the same convention here.

learning paradigm, in which we incorporate additional structural characteristics into the dictionary atoms to be identified, so that they be especially receptive to local deviations in the measured wavefields (and in turn, highly receptive to the wave propagation characteristics in the neighborhoods of local material anomalies). We motivate and describe this approach in the following section.

III. Dictionary Learning for Anomaly Detection

III.1. Dictionary learning for local feature extraction

The objective of anomaly detection is to discover regions that behave abnormally, i.e., whose behavior significantly deviates from that of their surroundings. In the realm of structural diagnostics, the term *anomaly* encompasses a wide variety of geometrical and material abnormalities, including damage zones, manufacturing defects and inclusions. In this discussion, we pursue the inference and triangulation of structural anomalies through the analysis of the structure’s dynamic response, which consists of a collection of displacement time histories acquired at a set of discrete points on the structure’s surface. Our contention is that it is possible to detect the signature of these anomalies by learning appropriate dictionaries of the response data and by decoding the spatial information contained in their data structure.

We assume that the regions containing physical anomalies are characterized by local perturbations of their acoustic properties (e.g. elastic moduli, density); these, in turn, induce localized features in the response, which are reflected, although possibly difficult to detect, in the kinematic time histories of the material points which lie inside the anomalous regions. If we invoke the interpretation of the atoms of a dictionary as pseudo modes (i.e., *spatial* descriptors of deformation), we can assume that the effect of localization would manifest as spiky regions in one or more of the atoms, which, from a data standpoint, would correspond to sparse structures in the columns of the dictionary. We conclude that, in order to equip a dictionary with the ability to detect anomalies, we need to formally enforce some kind of sparsity constraint on its atoms. On the other hand, the localized features associated with the anomalies coexist in the response with smoother fields describing the global (and dominant) behavior of the structure; therefore, a dictionary that properly captures an anomaly is unlikely to provide a sufficiently accurate representation of the response field as a whole, and viceversa. In order to reconcile this dichotomy, we propose a two-dictionary representation of the form

$$\mathbf{X} \approx \mathbf{D}_1 \mathbf{A}_1 + \mathbf{D}_2 \mathbf{A}_2, \quad (6)$$

as illustrated in Fig. 2. This can be alternatively interpreted as a decomposition of \mathbf{D} into two separate dictionaries, so that $\mathbf{D} = [\mathbf{D}_1 \mathbf{D}_2]$, where the $N \times K_1$ dictionary \mathbf{D}_1 is dedicated to capture the localized features of the response and the $N \times K_2$ dictionary \mathbf{D}_2 guarantees that the bulk response is sufficiently well approximated. Each dictionary is obtained through a dedicated optimization problem with appropriate constraints: specifically, for a user-specified constant $\Gamma > 0$, we impose on the atoms of \mathbf{D}_1 the sparsity-promoting constraint $\|(\mathbf{d}_1)_j\|_2^2 + \Gamma \|(\mathbf{d}_1)_j\|_1 \leq 1$ for each column $j = 1, 2, \dots, K_1$, while we retain the original column-wise constraints from the original dictionary learning formulation on the columns of \mathbf{D}_2 (that $\|(\mathbf{d}_2)_j\|_2^2 \leq 1$ for all $j = 1, \dots, K_2$). As we will see, the constraint on \mathbf{D}_2 tends to yield diffuse atoms, which *de facto* display spatial smoothness, while the additional ℓ_1 -based constraint promotes sparsity on the columns of \mathbf{D}_1 . Together, \mathbf{D}_1 and \mathbf{D}_2 form a representation that encompasses the dominant, smooth dynamic




Figure 2: Schematic illustration of the two-dictionary approximate representation: $\mathbf{X} \approx \mathbf{D}_1 \mathbf{A}_1 + \mathbf{D}_2 \mathbf{A}_2$ with $\mathbf{D}_1 \in \mathbb{R}^{N \times K_1}$, $\mathbf{D}_2 \in \mathbb{R}^{N \times K_2}$, $\mathbf{A}_1 \in \mathbb{R}^{K_1 \times T}$, $\mathbf{A}_2 \in \mathbb{R}^{K_2 \times T}$.

behavior as well as the spatially sparse signature of potential anomalies. The search for $\mathbf{D}_1, \mathbf{D}_2, \mathbf{A}_1, \mathbf{A}_2$ is done iteratively as detailed in Algorithm 1.

In the minimization problem of Eq. 9, the constraint on the individual columns of dictionary \mathbf{D}_1 are of the form $\|(\mathbf{d}_1)_j\|_2^2 + \Gamma \|(\mathbf{d}_1)_j\|_1 \leq 1$. Note that the Γ parameter effectively governs the number of non-zero terms in each atom. The larger Γ , the more stringent the sparsity constraint and the more zero elements in each atom, which in return makes the spatially sparse features more prevalent in the atoms of \mathbf{D}_1 . However, if Γ is too large, the constraint becomes excessively stringent and the cost function is trivially minimized with $\|\mathbf{D}_1\|_0 = 0$. On the other hand, if Γ is too small, the constraint is not sufficiently invoked and the atoms of \mathbf{D}_1 effectively lose the ability to highlight any localized features. The determination of the appropriate value of Γ is computed iteratively starting with a low initial guess Γ_0 and increasing it by an amount δ until the columns of \mathbf{D}_1 have a sufficiently low number ε of non-zero entries. In our implementation, the solution of the individual dictionary learning problems involved at each stage of iteration is carried out using the open-source sparse modeling software SPAMS (available at <http://spams-devel.gforge.inria.fr>).

Let us recall at this point that the role of the regularization parameter λ in the cost function of each dictionary learning problem is to control the trade-off between the accuracy of the approximation and the achievement of a parsimonious dictionary representation involving few atoms. In the present implementation, the selection of λ has been conducted following a heuristic approach such that the error in the approximation remained below acceptable bounds ($\approx 10\%$ of \mathbf{X}).

III.2. A super-atom approach to capture persistence and enhance anomaly detection

Important parameters of the algorithm are the user-defined parameters K_1 and K_2 , which represent the assumed number of atoms in \mathbf{D}_1 and \mathbf{D}_2 , respectively. If we let K_1 be large, we have the opportunity to examine a high number of sparse atoms. This richness of sparse descriptors can be exploited to better determine if the system contains an anomaly. It is, in fact, possible that some sparse atoms may not display the “true” physical anomalies, but rather other spurious short-wavelength features, e.g. cusps associated with boundary effects, which may manifest as localized features on the perimeter of the structure, or sharp artifacts due to noise. Nevertheless, we note that the true anomalies are a *persistent* feature in the sparse dictionary, i.e., features that are consistently observed across the set of sparse atoms and at consistent locations within each atom. In order to capture this *persistence* attribute, we propose a post-processing aggregation step designed to intelligently aggregate layers of data from multiple atoms in a way

Input: Data cube $\mathbf{X} \in \mathbb{R}^{N \times T}$
 Number of atoms in dictionaries: $K_1, K_2 > 0$
 Regularization parameters: $\lambda, \Gamma > 0$
 Sparsity parameter $\varepsilon > 0$
 Increment parameter $\delta > 0$
Output: Dictionaries \mathbf{D}_1 (sparse) and \mathbf{D}_2 (diffuse)

Initialize:

$$\begin{aligned} & \underset{\mathbf{D}_1 \in \mathbb{R}^{N \times K_1}, \mathbf{A}_1 \in \mathbb{R}^{K_1 \times T}}{\text{minimize}} && \sum_{t=1}^T \|\mathbf{x}_t - \mathbf{D}_1(\mathbf{a}_1)_t\|_2^2 + \lambda \|(\mathbf{a}_1)_t\|_1 \\ & \text{subject to} && \|(\mathbf{d}_1)_j\|_2^2 + \Gamma \|(\mathbf{d}_1)_j\|_1 \leq 1 \quad \forall \quad 1 \leq j \leq K_1 \end{aligned} \quad (7)$$

while $\|\mathbf{D}_1\|_0 \geq \varepsilon$ **do**

$\Gamma = \Gamma + \delta$
 $\mathbf{X}' = \mathbf{X} - \mathbf{D}_1 \mathbf{A}_1$

$$\begin{aligned} & \underset{\mathbf{D}_2 \in \mathbb{R}^{N \times K_2}, \mathbf{A}_2 \in \mathbb{R}^{K_2 \times T}}{\text{minimize}} && \sum_{t=1}^T \|\mathbf{x}'_t - \mathbf{D}_2(\mathbf{a}_2)_t\|_2^2 + \lambda \|(\mathbf{a}_2)_t\|_1 \\ & \text{subject to} && \|(\mathbf{d}_2)_j\|_2^2 \leq 1 \quad \forall \quad 1 \leq j \leq K_2 \end{aligned} \quad (8)$$

$\tilde{\mathbf{X}} = \mathbf{X} - \mathbf{D}_2 \mathbf{A}_2$

$$\begin{aligned} & \underset{\mathbf{D}_1 \in \mathbb{R}^{N \times K_1}, \mathbf{A}_1 \in \mathbb{R}^{K_1 \times T}}{\text{minimize}} && \sum_{t=1}^T \|\tilde{\mathbf{x}}_t - \mathbf{D}_1(\mathbf{a}_1)_t\|_2^2 + \lambda \|(\mathbf{a}_1)_t\|_1 \\ & \text{subject to} && \|(\mathbf{d}_1)_j\|_2^2 + \Gamma \|(\mathbf{d}_1)_j\|_1 \leq 1 \quad \forall \quad 1 \leq j \leq K_1 \end{aligned} \quad (9)$$

end

return $\mathbf{D}_1, \mathbf{D}_2$

Algorithm 1: Iterative procedure to learn two dictionaries (*sparse* and *diffuse*) from the response data.

that emphasizes the most persistent features. The result of this step is a kind of *super-atom* that highlights spatial locations where persistent activity is present across a significant number of the sparse dictionary atoms. The construction of the super-atom proceeds as follows. We consider a partition of the domain into $M_1 \times M_2$ rectangular (identical in size and shape) regions. Since the atoms spatially span the entire domain, they are all partitioned in similar fashion, such that \mathbf{d}_j^k denotes the k^{th} partition of the \mathbf{d}_j atom. For each partition, we sweep the atoms of the sparse dictionary and we check if a feature is consistently observed in that partition across the set, by counting how many atoms contain at least one non zero entry inside the selected partition. If this number is sufficiently large, we aggregate local contributions from all the atoms in the dictionary to form the corresponding partition of the super-atom. Note that this criterion weighs (possibly relatively mild) contributions that are observed over a large number of atoms more heavily than others that may be prominent (amplitude wise), but are observed only in a few atoms. This reflects the notion that the signature of physical anomalies is often elusive but persistent across the dictionary, while spurious sharp features, which can dominate the response amplitude wise, are inconsistently detected across the dictionary. The construction of the super-atom (summarized in Algorithm 2) features two parts: the

Input: Sparse Dictionary $\mathbf{D}_1 = [\mathbf{d}_1 \ \mathbf{d}_2 \ \dots \ \mathbf{d}_{K_1}] \in \mathbb{R}^{N \times K_1}$

Domain partition sizes: $M_1, M_2 > 0$

Sparsity thresholds $\varepsilon, \delta > 0$

Output: “Super-atom” \mathbf{S} (with vectorized partitions $\{\mathbf{S}^k\}$)

Initialize: $M = M_1 M_2$; $Q = N/M$

$\mathbf{S}^k = \mathbf{0} \in \mathbb{R}^Q$ for $k = 1, \dots, M$

Partition: Divide each \mathbf{d}_j , $j = 1, \dots, K_1$, into length Q sub vectors $\{\mathbf{d}_j^k\}_{k=1}^M$

```

for  $k = 1$  to  $M$  do
    if  $\sum_{j=1}^{K_1} \mathbf{1}_{\{\|\mathbf{d}_j^k\|_0 > 0\}} \geq \varepsilon$  then
         $\mathbf{S}^k \leftarrow \sum_{j=1}^{K_1} \|\mathbf{d}_j^k\|_0$ 
    end
    for  $i = 1$  to  $Q$  do
        if  $|S_i^k| \leq \delta$  then
             $S_i^k \leftarrow 0$ 
        end
    end
end
return  $\{\mathbf{S}^k\}$ 

```

Algorithm 2: Construction of super-atom aggregating contributions from different sets of sparse atoms according to criteria of persistence. The function $\mathbf{1}_{\{\cdot\}}$ denotes the *indicator function*, which takes the value 1 when the event specified in the subscript is true, and 0 otherwise.

first implements the atom aggregation procedure; the second performs a search over the identified partitions, with the objective of automatically identifying, through amplitude thresholding, the partition containing the anomalies. This last step is meant to forsake the need for visual inspection of the super-atom and *de facto* makes the anomaly identification fully automatic.

Note that there may be instances where the algorithm does not find any partition that contains any sufficiently persistent localized features; this scenario corresponds to the case of a pristine structure. The ability to avoid false positives in anomaly-free scenarios is a major strength of this methodology, in contrast with methods that do not explicitly control for false positives, e.g. our own previous work [42], (which utilized a simpler, principal components analysis-based anomaly detection and localization method).

The benefits of a super-atom representation are fully realized when we consider problems in which the identification of an anomaly is impossible, or ambiguous, even through the prism of the sparse atom dictionary. This scenario is encountered when the signature of the anomaly is very small, due, for example, to a critically small size of the defect or to competition (in terms of sparsity) from spurious response features, or in dealing with noisy data. In short, the signature of an anomaly may be minute, and hence not visually detectable, in the individual atoms, but can become very clear through its aggregation over many atoms.

It is worthwhile to note that our super-atom post-processing method is a bit of a departure from existing dictionary learning approaches, which typically utilize directly the atoms identified by the learning procedure(s) without further refinement. Our motivation for adopting this additional step is twofold. First,

as described above, the approach seeks to identify spatially persistent features in the response data by aggregating (in a nonlinear manner) features identified in the atoms of the sparse dictionary. In addition, we note that while dictionary learning problems are easy to motivate and specify, their highly non-convex nature makes their exact numerical solution computationally challenging. In practice, existing algorithmic approaches rely either on convex relaxation and alternating optimization or greedy methods and can only be guaranteed to converge to local minima of the corresponding objective function. In this sense, our post processing step may be viewed as an augmentation designed to glean additional information from the computational solution of the dictionary learning problem.


IV. Local Feature Extraction in a Transient Wavefield

In this section we provide several numerical demonstrations of the proposed anomaly identification method, on both simulated data and experimental data obtained using a scanning laser Doppler vibrometer.


IV.1. Results from numerical simulations

First we test the approach against the problem of circular-crested transversal waves excited in a thin plate by an out-of-plane point force applied at one node. We consider flexural waves modeled according to Mindlin’s plate theory. The choice of flexural waves is here primarily motivated by the simplicity of the corresponding finite element simulation and the inherent simplifications resulting from having a single-mode wave solution. Nevertheless, since the method is based on the elaboration of the spatial patterns of time-evolving wavefields, without invoking any specific physical characteristics of the waves, the analysis would hold for other types of out-of-plane waves, such as Lamb waves and Rayleigh waves, and even for longitudinal and shear waves in thin structures exhibiting some in-plane deformation on the structure’s surface. We use the finite element method (FEM) to construct the stiffness and mass matrices of the system and employ a time marching scheme to simulate the propagating wavefield. Our virtual specimen is a rectangular, thin Aluminum sheet with dimensions $L_x = 1.0\text{m}$, $L_y = 0.5\text{m}$ and thickness $h = 0.002\text{m}$ and the following material parameters: Young’s modulus $E = 71\text{GPa}$, Poisson ratio $\nu = 0.33$, density $\rho = 2700\text{Kg/m}^3$. The domain is discretized with a structured mesh comprising 400×200 square elements, which is verified to be sufficiently fine to avoid spurious numerics-induced dispersive effects or artificial noise in the data. The excitation frequency is a 5-cycle tone burst with carrier frequency $f_c = 100\text{KHz}$, which induces a wavefield with wavelength λ such that $L_x/\lambda \approx 30$. At the end of simulation, our data cube \mathbf{X} consists of a series of snapshots of the propagating wavefield at the selected T time instants.

The anomalies are introduced in the model by relaxing the Young’s modulus of the material by two orders of magnitude within small regions of the domain. This can model the effects of partial holes, soft inclusions or localized regions with degraded material properties. The anomalies behave as scatterers, which act as localized sources within the domain that are triggered after some delay with respect to the applied excitation. As the point source of excitation is itself a localized feature, we expect it to pose some ambiguity for the sparse coding algorithm. In order to (partially) filter out this effect and discriminate between the “true” anomalies and the excitation, we truncate the first 25% time instants of the response (corresponding to the early stages of propagation), in which we expect the response to be dominated by the excitation as the wavefield is localized in the neighborhood of the excitation point. For the same reason, a $0.01L_x$ -thick layer immediately close to the boundary where the excitation is applied, is *a priori* excluded from the analysis.



(a) Diffuse atoms (9 random samples of 16-atom dictionary)



(b) Sparse atoms (9 random samples of 100-atom dictionary)

Figure 3: Samples of atoms from diffuse and sparse dictionaries for a numerically generated wavefield with point anomaly. Arrows are used to assist the visualization of the spikes corresponding to the anomaly.

In Fig. 3 we show a sample of 9 atoms of a low-dimensional ($K_2 = 16$) diffuse dictionary \mathbf{D}_2 and a sample of 9 randomly selected atoms of a higher-dimensional ($K_1 = 100$) sparse dictionary \mathbf{D}_1 for a case with a single point anomaly of coordinates $(0.25L_x, 0.33L_y)$. The diffuse atoms (Fig. 3a) essentially capture characteristic snapshots of the wavefield at different instants of simulation. In contrast, the sparse dictionary (Fig. 3b) captures a spatially rarefied representation of the wavefield, in which several atoms display a distinguishable and isolated signature of the anomaly in the form of a localized spike in their topology; this signature is, however, quite elusive, as the anomalous feature is not ubiquitously observed across the entire set and the task of discriminating it from other speckles in each atom is prohibitive. As expected, the inference benefits

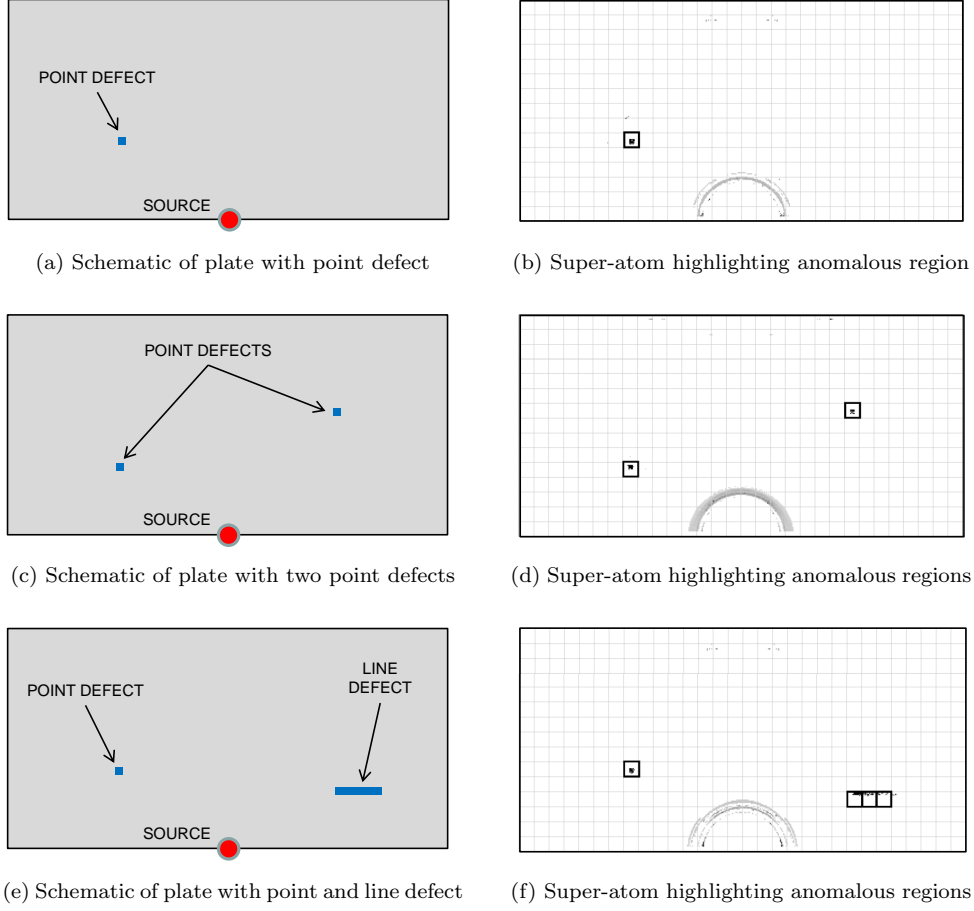


Figure 4: Anomaly detection and triangulation through the prism of super-atom representation. Point and line anomalies are successfully extracted from the wavefields and triangulated.

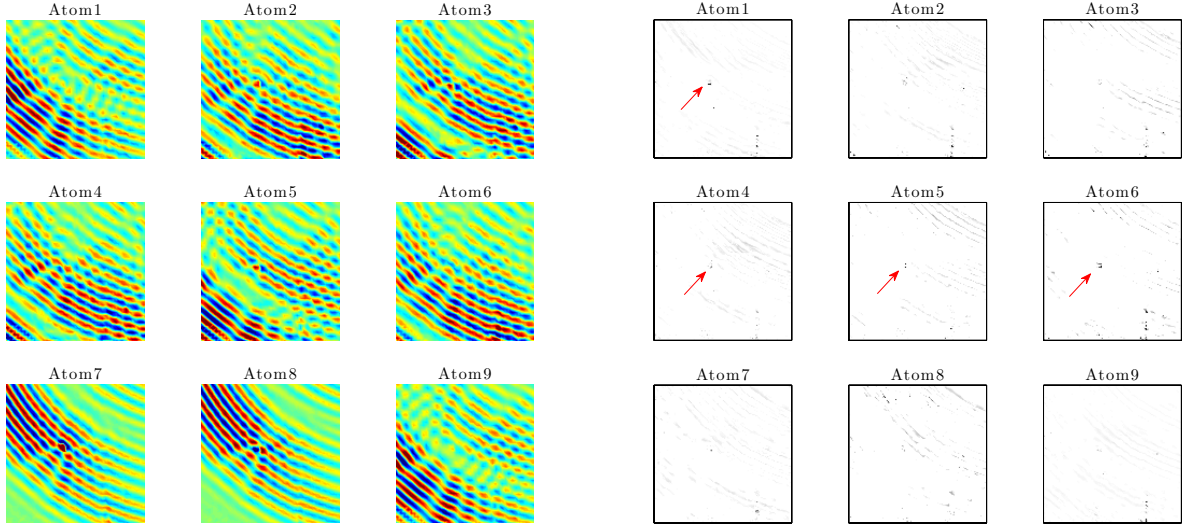
vastly from a super-atom representation (Fig. 4b), in which the sparse features are weighted according to their persistence across the entire sparse dictionary. The sporadically occurring features that contaminate the sparse atoms are now filtered out and we are able to clearly pinpoint the anomaly, as visible from the comparison with the schematic of Fig. 4a. We note that the anomaly can even be triangulated without visual inspection as its host partition is automatically identified by the algorithm (and highlighted in the figure by a thicker border). In the remaining portion of Fig. 4 we further explore the performance of the method against cases with more challenging anomaly landscapes. In Figs. 4c and 4d we show the super-atom performance for a plate with two defects, where both anomalies are correctly detected and triangulated. The final case in Figs. 4e and 4f tests whether a crack, which in this case can be thought of as a contiguous collection of anomalies, can be identified even in the presence of another scatterer. The crack is simulated by reducing the Young's modulus and density of the material by several orders of magnitude inside a two-element thick horizontal layer extending between points $(0.74L_x, 0.20L_y)$ and $(0.84L_x, 0.20L_y)$. The signature of the crack is well highlighted in the super-atom; if we relax the search criterion and we let the algorithm find the 4 most anomalous partitions, we are able to automatically identify the entire length of the defect.

IV.2. Results from experimental data

We now test the efficacy of the dictionary learning algorithm against experimental data through the benchmark problem of a thin plate with a localized defect, the physical analogue of the numerical case presented in the previous subsection. Two cases, using a similar specimen but with considerably different types of anomalies, are presented. Here we consider a square Aluminum sheet of dimensions $L_x = L_y = 0.61\text{m}$ and thickness $h = 0.0394\text{m}$ and we apply a tone burst excitation with carrier frequency $f_c = 200\text{kHz}$. The induced wavefield is reconstructed from surface velocity data via a Scanning Laser Doppler Vibrometer (SLDV), the Polytec PSV-400-3D. Three scanning heads are used to shoot three independent laser beams onto each node of a predefined scanning grid, which allows capturing the out-of-plane and in-plane velocity components of the points on the plate’s surface; in this work, we limit our analysis to the out-of-plane velocity component. The plate surface is discretized with an approximately square scanning grid consisting of 120×120 scan points. The grid is deliberately chosen to be very fine to enable resolution of minute features in the wavefield; experiments with coarser grids can be easily performed *a posteriori* by subsampling the data in space. To make the wavefield data more amenable to region partitioning, we apply linear interpolation to fit the original grid data to a Cartesian grid. Since surface roughness can cause signal dropouts, or temporary severe signal attenuation, built-in signal enhancement and speckle tracking features are enabled during the scan. Moreover we average 75 realizations at each scan point to filter optical and/or mechanical noise (signal stacking). Despite many precautions taken during the scan, the acquired signals are tainted by some mechanical and optical noise; the data is therefore cleaned by post-processing the time history signals with a band pass filter with a 200 kHz center frequency and a bandwidth of 100 kHz. Time spacing between acquisitions at different points is enforced by a repetition frequency of 9Hz, which allows for over 1000 boundary reflections between successive acquisitions and effectively allows for full attenuation of the wave before a new acquisition is made.


In the first experimental case we consider a localized anomaly introduced in the form of a circular, flat-bottomed cylindrical not-through hole, shown in detail in the schematic of Fig. 6a. The hole is drilled on the back surface of the plate to a depth of approximately 0.0295 m - corresponding to a 75% reduction in local thickness - and is expected to behave as a strong scatterer. Nevertheless, it is worth pointing out that the scan is performed over a portion of the *pristine* front face of the plate (also highlighted in Fig. 6a), which shows no visual evidence of the defect; the experiment mimics diagnostics conditions encountered in the inspection of certain thin-walled structures - such as mounted aircraft panels, pipes, etc. - for which it may be prohibitive to have direct optical access to the interior surfaces. In Fig. 5 we show the diffuse atoms and a randomly selected sub-sample of 100 computed sparse atoms. Similar to the numerical case, the diffuse atoms (Fig. 5a) effectively display snapshots of the wavefield at different time instants of the experiment, while the sparse atoms (Fig. 5b) distill the sparse, or localized features of the response. While the presence of a distinguishable localized feature is observed in several of the atoms at the location of the anomaly, the inference is again contaminated by other sparse features, which introduces an element of ambiguity. Through the lens of the super-atom, as seen in Fig. 6b, the “true” anomaly is clearly displayed and highlighted.

Our second experimental case is an intriguing testimony to the method’s ability to detect subtle changes induced in a wavefield by minute or superficial defects. The almost serendipitous way in which this detection was obtained speaks volumes about the agnostic capabilities of the method. A plate with identical specifi-




(a) Diffuse atoms (9 random samples of 16-atom dictionary) (b) Sparse atoms (9 random samples of 100-atom dictionary)

Figure 5: Samples of atoms from diffuse and sparse dictionaries for a laser-acquired wavefield with strong scatterer. Arrows are used to assist the visualization of the spikes corresponding to the anomaly.



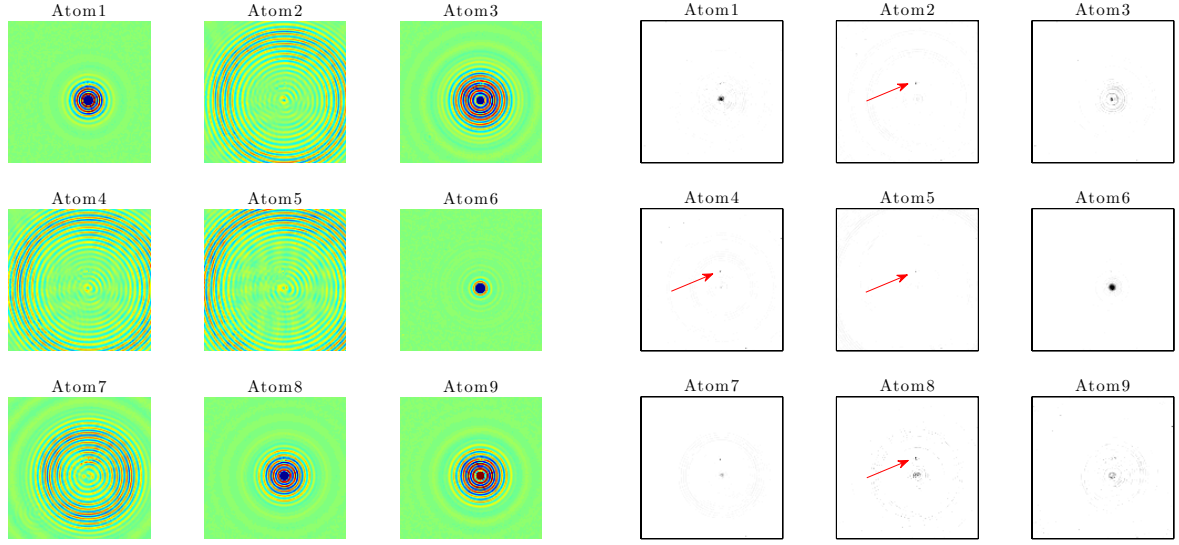
(a) Schematic of plate with highlighted anomaly (picture detail) and scanned area



(b) Super-atom of the scanned area successfully identifying the anomaly


Figure 6: Anomaly detection and triangulation through the prism of super-atom representation. The defect corresponding to the hole is successfully triangulated.

cations as above (without the drilled hole), unknowingly possessed a very thin, approximately 0.1mm-thick deposit of glue - remnant of an actuator that had been previously glued to its back surface (see picture detail in the schematic of Fig. 8a). A preliminary scan was performed for the purpose of calibrating the equipment and obtaining a preliminary characterization of the wavefield. Since we were ostensibly dealing




(a) Diffuse atoms (9 random samples of 16-atom dictionary) (b) Sparse atoms (9 random samples of 100-atom dictionary)

Figure 7: Samples of atoms from diffuse and sparse dictionaries for a laser-acquired wavefield with weak scatterer. Arrows are used to assist the visualization of the spikes corresponding to the anomaly.



(a) Schematic of plate with highlighted anomaly and scanned area



(b) Superatom of the scanned area identifying source and anomaly

Figure 8: Anomaly detection and triangulation through the prism of super-atom representation. The defect corresponding to the glue residual is successfully triangulated.

with a pristine plate, we expected a radially symmetric, unperturbed wavefield. Even though some marginal distortion potentially ascribable to a weak scatterer was detected upon visual inspection, the observed wavefield initially confirmed the prediction of an essentially pristine structure. However, when the data was fed

to the sparse coding algorithm, a number of localized features were displayed in the sparse atoms (Fig. 7b). One of them corresponds to the excitation point, which here lies in the middle of the domain - a region that is not *a priori* excluded from the analysis; this is not surprising, as excitation sources are inherently spatially localized features. However, we also noted that many of the sparse atoms contained a second feature above the point of excitation. This was emphasized in the super-atom representation (Fig. 8b), which indeed suggests the presence of a (weak) scatterer. Further inspection of the plate back surface revealed the presence of the aforementioned defect at the location indicated by the super-atom.

V. Concluding Remarks and Future Work

In this work we have introduced a methodology to detect and triangulate anomalies in the response of solid media based on criteria of spatial sparsity and enabled by sparse coding algorithms. The method represents, to the best of our knowledge, the first attempt to use structurally-tuned dictionary learning algorithms in the context of structural diagnostics and, in general, one of the few existing efforts aimed at performing anomaly detection from spatially reconstructed wavefields using techniques of computer vision and image processing. We have shown that the method can identify spatially localized anomalies in the data fields through a decomposition of the response in two sets of pseudo modes: diffuse ones, which capture the smooth part of the response, and sparse ones, which contain the signature of the anomalies. The sparse coding algorithms have been complemented with the assembly of super-atoms that intelligently aggregate the information from the sparse dictionaries to filter the localized features that are due to physical anomalies from those due to noise or other competing boundary-induced or numerics-induced mechanisms.

The method is crowned by a postprocessing feature that allows a convenient virtual decomposition of the domain in rectangular partitions for automatic identification of the regions containing the anomalies. The benefits of an automatic interpretation that bypasses the need for direct visual observations of the wavefield (or of the atoms of its dictionaries) are felt in the context of possible *multi-step* sampling and detection procedures, where the inference would be made in several stages conducted over nested sub-domains, and with increased accuracy, to iteratively identify smaller and smaller subsets of the material domain that may contain anomalies. This sort of adaptive “coarse-to-fine” sampling strategy would enable agile and fast sensing and detection procedures which could in return enhance the applicability and competitiveness of image-processing-based diagnostics methods. This is the objective of current investigations; an account of this is left for future work.

Another interesting avenue for future work consists of testing the approach in the context of heterogeneous materials, in which the material properties could potentially feature large deviations from an ideal case even far from the to-be-detected anomalies. This type of scenario would ultimately testify to the agnostic properties of the approach, as in those cases the ability to triangulate anomalies without precognition of the mechanical properties of the medium would fully manifest.

References

- [1] W.J. Staszewski, C. Boller, and G. Tomlinson. *Health monitoring of aerospace structures: Smart sensors and signal processing*. Wiley & Sons, 2004.
- [2] J.L. Rose. A baseline and vision of ultrasonic guided wave inspection potential. *Journal of Pressure Vessel Technology*, 124(3):273–282, 2002.

- [3] J.-B. Ihn and F.-K. Chang. Pitch-catch active sensing methods in structural health monitoring for aircraft structures. *Structural Health Monitoring*, 7(1):5–19, 2008.
- [4] E. B. Flynn, M. D. Todd, P. D. Wilcox, B. W. Drinkwater, and A. J. Croxford. Maximum-likelihood estimation of damage location in guided-wave structural health monitoring. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 467(2133):2575–2596, 2011.
- [5] T. E. Michaels, J. E. Michaels, B. Mi, and M. Ruzzene. Damage detection in plate structures using sparse ultrasonic transducer arrays and acoustic wavefield imaging. In *AIP Conference Proceedings*, volume 760, page 938, 2005.
- [6] L. Liu, S. Liu, and F.-G. Yuan. Damage localization using a power-efficient distributed on-board signal processing algorithm in a wireless sensor network. *Smart Materials and Structures*, 21(2):025005, 2012.
- [7] Y. Lu, X. Wang, J. Tang, and Y. Ding. Damage detection using piezoelectric transducers and the Lamb wave approach: II. Robust and quantitative decision making. *Smart Materials and Structures*, 17(2):025034, 2008.
- [8] Q. Wang and S. Yuan. Baseline-free imaging method based on new PZT sensor arrangements. *Journal of Intelligent Material Systems and Structures*, 20(14):1663–1673, 2009.
- [9] C. Prada and M. Fink. Eigenmodes of the time reversal operator: A solution to selective focusing in multiple-target media. *Wave Motion*, 20(2):151–163, 1994.
- [10] F. Foroozan and A. Asif. Time reversal based active array source localization. *IEEE Transactions on Signal Processing*, 59(6):2655–2668, 2011.
- [11] S.S. Kessler, S.M. Spearing, and C. Soutis. Damage detection in composite materials using lamb wave methods. *Smart Materials and Structures*, 11(2):269, 2002.
- [12] G.R. Kirikera, O. Balogun, and S. Krishnaswamy. Adaptive Fiber Bragg Grating Sensor Network for Structural Health Monitoring: Applications to Impact Monitoring. *Structural Health Monitoring*, 10(1):5–16, 2011.
- [13] C. Farrar and K. Worden. *Structural health monitoring: A machine learning perspective*. Wiley & Sons, 2013.
- [14] S. Zhongqing and L. Ye. Lamb wave-based quantitative identification of delamination in CF/EP composite structures using artificial neural algorithm. *Composite Structures*, 66(1):627 – 637, 2004.
- [15] Y. Lu and J. E. Michaels. Ultrasonic Signal Decomposition via Matching Pursuit with an Adaptive and Interpolated Dictionary. In D. O. Thompson and D. E. Chimenti, editors, *Review of Progress in Quantitative Nondestructive Evaluation*, volume 894 of *American Institute of Physics Conference Series*, pages 579–586, 2007.
- [16] S. Das, A. Papandreou-Suppappola, X. Zhou, and A. Chattopadhyay. On the use of the matching pursuit decomposition signal processing technique for structural health monitoring. *Smart Structures and Materials 2005: Smart Structures and Integrated Systems*, 5764(1):583–594, 2005.
- [17] S. Das, I. Kyriakides, A. Chattopadhyay, and A. Papandreou-Suppappola. Monte carlo matching pursuit decomposition method for damage quantification in composite structures. *Journal of Intelligent Material Systems and Structures*, 20(6): 647–658, 2009.
- [18] S.G. Mallat and Z. Zhang. Matching pursuits with time-frequency dictionaries. *IEEE Trans Signal Processing*, 41(12): 3397 – 3415, 1993.
- [19] S. Das, A. N. Srivastava, and A. Chattopadhyay. Classification of damage signatures in composite plates using one-class SVMs. In *Proc. IEEE Aerospace Conference*, pages 1 – 19, 2007.
- [20] A. Azimipana and S. ShahbazPanahi. Experimental results of compressive sensing based imaging in ultrasonic non-destructive testing. In *2013 IEEE 5th International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, pages 336–339, 2013.
- [21] V. Sharma, S. Hanagud, and M. Ruzzene. Damage index estimation in beams and plates using laser vibrometry. *AIAA Journal*, 44(4):919–923, 2006.
- [22] T.E. Michaels, J.E. Michaels, and M. Ruzzene. Frequency-wavenumber domain analysis of guided wavefields. *Ultrasonics*, 51(4):452 – 466, 2011.
- [23] R. Basri and W. K. Chiu. Numerical analysis on the interaction of guided Lamb waves with a local elastic stiffness reduction in quasi-isotropic composite plate structures. *Composite Structures*, 66:87–99, 2004.
- [24] D. Alleyne and P. Cawley. A two-dimensional Fourier transform method for the measurement of propagating multimode signals. *Journal of the Acoustical Society of America*, 89:1159–1168, 1991.
- [25] M. Ruzzene. Frequency-wavenumber domain filtering for improved damage visualization. *Smart Materials and Structures*, 16(6):2116, 2007.
- [26] H. Sohn, D. Dutta, J.Y. Yang, M. DeSimio, S. Olson, and E. Swenson. Automated detection of delamination and disbond from wavefield images obtained using a scanning laser vibrometer. *Smart Materials and Structures*, 20(4):045017, 2011.

- [27] Y.-K. An, B. Park, and H. Sohn. Complete noncontact laser ultrasonic imaging for automated crack visualization in a plate. *Smart Materials and Structures*, 22(2):025022, 2013.
- [28] L. Itti, C. Koch, and E. Niebur. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(11), 1998.
- [29] L. Itti and C. Koch. Computational modelling of visual attention. *Nature*, 2, March 2001.
- [30] J. Yan, M. Zhu, H. Liu, and Y. Liu. Visual saliency detection via sparsity pursuit. *IEEE Signal Proc. Letters*, 17(8), 2010.
- [31] X. Shen and Y. Wu. A unified approach to salient object detection via low rank matrix recovery. In *Proc. Computer Vision and Pattern Recognition*, 2012.
- [32] A. Patcha and J. Park. An overview of anomaly detection techniques: Existing solutions and latest technological trends. *Computer Networks*, 51(12), 2007.
- [33] V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 2009.
- [34] S. S. Chen, D. L. Donoho, and M. A. Saunders. Atomic decomposition by basis pursuit. *SIAM Rev.*, 43(1):129–159, 2001.
- [35] E. Candes, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2):489–509, February 2006.
- [36] D. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, April 2006.
- [37] E. J. Candès and T. Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on Information Theory*, 52(12):5406–5425, December 2006.
- [38] J. Haupt and R. Nowak. Signal reconstruction from noisy random projections. *IEEE Transactions on Information Theory*, 52(9):4036–4048, September 2006.
- [39] J. A. Tropp. Greed is good: Algorithmic results for sparse approximation. *IEEE Transactions on Information Theory*, 50(10):2231–2242, 2004.
- [40] J. A. Tropp. Just relax: Convex programming methods for identifying sparse signals in noise. *IEEE Transactions on Information Theory*, 52(3):1030–1051, 2006.
- [41] A.M. Bruckstein, D.L. Donoho, and M. Elad. From sparse solutions of systems of equations to sparse modeling of signals and images. *SIAM Rev.*, 51(1):34–81, 2009.
- [42] S. Gonella and J.D. Haupt. Automated defect localization via low rank plus outlier modeling of propagating wavefield data. *IEEE Transactions on Ultrasonics, Ferroelectrics and Frequency Control*, 60(12):2553–2565, 2013.
- [43] B. A. Olshausen and D. J. Field. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision Research*, 37:3311–3325, 1997.
- [44] M. S. Lewicki, T. J. Sejnowski, and H. Hughes. Learning overcomplete representations. *Neural Computation*, 12:337–365, 1998.
- [45] K. Kreutz-Delgado, J. F. Murray, B. D. Rao, K. Engan, T.-W. Lee, and T. J. Sejnowski. Dictionary learning algorithms for sparse representation. *Neural computation*, 15(2):349–396, 2003.
- [46] M. Aharon, M. Elad, and A. Bruckstein. K-SVD: An algorithm for designing overcomplete dictionaries for sparse representation. *IEEE Transactions on Signal Processing*, 54(11):4311–4322, 2006.
- [47] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online learning for matrix factorization and sparse coding. *J. Mach. Learn. Res.*, 11:19–60, 2010.
- [48] J. Mairal, M. Elad, and G. Sapiro. Sparse representation for color image restoration. *IEEE Transactions on Image Processing*, 17(1):53–69, 2008.
- [49] G. Peyré. Sparse modeling of textures. *Journal of Mathematical Imaging and Vision*, 34(1):17–31, 2009.
- [50] I. Ramirez, P. Sprechmann, and G. Sapiro. Classification and clustering via dictionary learning with structured incoherence and shared features. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*, pages 3501–3508, 2010.
- [51] S. Zubair, F. Yan, and W. Wang. Dictionary learning based sparse coefficients for audio classification with max and average pooling. *Digital Signal Processing*, 23(3):960 – 970, 2013.
- [52] S. Rambhatla and J. Haupt. Semi-blind source separation via sparse representations and online dictionary learning. In *Proc. Asilomar Conf. on Signals, Systems, and Computers*, Pacific Grove, CA, November 2013.
- [53] Y. Li. Dictionary learning based multitask image restoration. In *Proc. Intl. Conf. on Image and Signal Processing*, pages 364–368, 2012.
- [54] J. Mairal, F. Bach, J. Ponce, and G. Sapiro. Online dictionary learning for sparse coding. In *Proc. ICML*, 2009.